arXiv:1603.04387v1 [cs.NI] 14 Mar 2016

NetMemex: Providing Full-Fidelity Traffic Archival

Hyeontaek Lim, Vyas Sekar, Yoshihisa Abe, David G. Andersen
Carnegie Mellon University

ABSTRACT

NetMemex explores efficient network traffic archival without any loss of information. Unlike NetFlow-like aggregation, NetMemex allows retrieving the entire packet data including full payload, which makes it useful in forensic analysis, networked and distributed system research, and network administration. Different from packet trace dumps, NetMemex performs sophisticated data compression for small storage space use and optimizes the data layout for fast query processing. NetMemex takes advantage of high-speed random access of flash drives and inexpensive storage space of hard disk drives. These efforts lead to a cost-effective yet high-performance full traffic archival system. We demonstrate that NetMemex can record full-fidelity traffic at near-Gbps rates using a single commodity machine, handling common queries at up to 90.1 K queries/second, at a low storage cost comparable to conventional hard disk-only traffic archival solutions.

1. INTRODUCTION

Archiving network traffic without loss of information is a useful foundation for activities including forensic analysis [12, 18, 24, 26, 29, 34], scientific research, particularly about networking and distributed systems [4, 7, 35], network administration [13, 34], and others.

However, the task of traffic archival is difficult, as recent applications impose technical challenges:

- **Data volume**. Applications often require recording a few or more days of full traffic for analysis (e.g., forensic analysis) [24, 30], resulting in a large amount of data to store (10.8+ TB/Gbps/day without compression).
- **Recording performance**. It has become more common to record high-speed traffic (e.g., 1 Gbps rate) [6, 11, 24, 29].
- **Query performance**. There is an increasing demand for fast query processing [14, 15, 24].

Previous solutions fail to fully address these challenges:

(1) *Lossy traffic archival*. There are numerous works that mitigate high storage consumption by aggregating or selectively recording traffic data [8, 24, 30]. However, this lossy approach is inadequate for applications that cannot tolerate any loss of information from the original traffic. A few compression schemes have been proposed for compressing traffic information [16], but they focus only on well-structured formats such as NetFlow records.

(2) *Imbalance between recording and query performance, storage use*. Many traffic archival systems have difficulty in achieving both high recording and query processing performance. One of the most widespread methods is to focus on recording throughput by viewing hard disk drives as a huge circular queue and streaming packets to the queue as they arrive [12, 13, 15, 24]. In this technique, query processing can waste a large amount of I/O bandwidth due to high false positive rates incurred by inefficient indexing over such an unstructured data layout. Some systems aim instead for high query performance [18], but they do not target high recording speed, so these approaches are only viable for offline or semi-online recording. These works do not perform detailed evaluation on the total storage cost, including the traffic data and any other metadata such as indexes, unless their system design focuses on reducing the storage space use [24, 30].

In this paper, we present NetMemex, a full-fidelity traffic archival system that simultaneously attains high recording and query performance and efficient storage. A single-node server-class system running NetMemex can archive Internet traffic of 0.62–1.15 Gbps, and it can handle up to 90.1 K common queries per second. While NetMemex employs flash drives, which are expensive relative to hard disk drives, because of NetMemex's data compression techniques, its total storage cost is similar to conventional hard disk only solutions.

The three key elements of our approach are:

- **Flow-oriented data reorganization and indexing** increases the efficiency of query processing and the effectiveness of data compression by transforming the packet stream into a flow-oriented form;
- **Employing both flash and hard disk** improves query latency by storing frequently and randomly accessed small data on flash while keeping less frequently accessed bulk data on hard disk; and
- **Extensive data compression** reduces the total storage cost by using effective data compression techniques, which are applied to the packet header and payload separately.

These approaches are not a combination of piecewise tech-

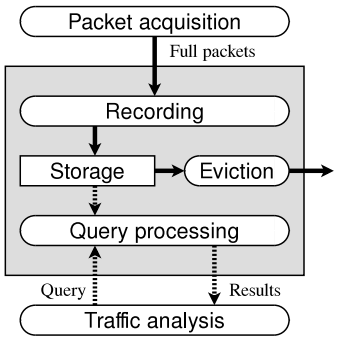

Figure 1: High-level view of a traffic archiving system (within the gray box) and its interaction with other system components (outside of the gray box). Solid arrows show data processing, storage, and eviction, and dashed arrows are data retrieval.

niques; they work together to help NetMemex achieve its performance and cost goals. For example, flow reordering substantially increases the effectiveness of compression, and data compression enables storing frequently accessed data on flash without a significant increase in the total storage cost.

2. BACKGROUND

We first describe a general traffic archiving system and its challenges, and then we show several data compression techniques and storage technologies that we use in NetMemex.

2.1 Traffic Archival Systems

A traffic archival system is a storage system specialized for network traffic data storage and retrieval. This system stores network packets and returns the stored information upon request.

Figure 1 shows typical components of a traffic archiving system. A traffic archiving system has a simple interface. The system receives captured packets from an external packet acquisition method, such as tcpdump/libpcap [37] and Endace DAG [13], and keeps them for a specific period of time. Stored packets and traffic information can be requested by traffic analysis software.

Inside the system, incoming packets go through the stages of recording and indexing, storage, and eviction. The recording and indexing phase accepts the input packets and stores them to storage. The packets can be transformed in this stage if there is a better data format available. As the storage fills, the system evicts old data to make room for new packets. The evicted data can be either discarded or stored in other storage for later analysis and longer archival; the auxiliary storage can be a larger-scale traffic archiving system or a large-scale ordinary storage cluster.

2.2 Data Compression

Dictionary-based compression techniques are most commonly used in everyday data compression. They maintain recently seen data patterns in a data structure called a dictionary. They use the dictionary to find a duplicate pattern in the new data, and they replace the pattern with a reference to the pattern in the dictionary. Since the reference can be encoded in smaller space than the original pattern in many cases, they can reduce the total space required to represent the original data.

A larger dictionary often (but not always) allows a higher compression ratio; unfortunately, it also typically decreases compression speed due to more expensive dictionary search operation. Therefore, fast compression techniques use small dictionary sizes (e.g., 32 KiB for zlib [17]). Redundancy elimination [3, 35], a form of dictionary-based compression on network routers, also maintain a small time window to operate at a high rate.

Data deduplication techniques [21], unlike dictionary-based compression, efficiently detect and reduce data redundancy across longer ranges. They use "fingerprints" to divide data stream into smaller blocks called *chunks*. Each chunk is hashed using a strong hash function (e.g., SHA-1) to generate a unique identifier, and they store only one copy of chunks with the same hash. Because it is possible to maintain an index of chunk hashes that are collected for a large amount of data (e.g., TBs) and a long period of time (e.g., days), data deduplication can save space even if redundant data are scattered across the entire data.

Dictionary-based compression and data deduplication are often used together to increase the compression effectiveness; each chunk is compressed using a dictionary-based compression technique while chunks are deduplicated using a data deduplication technique.

Header compression techniques [10, 19, 21] reduce the size of packet headers by exploiting redundancy in packets in the same flow (i.e., between packets sharing the same 5-tuple). For example, the sequence and acknowledge numbers in a TCP flow have small increments, which can be encoded in one or two octets instead of the full four octets. Header compression is typically applied between network links or end-points to reduce the bandwidth required to transit packets.

2.3 Storage Systems and Technologies

Log-structured file systems [20, 33] convert random writes to sequential writes by appending new data to a log. This approach takes advantage of the fact that many storage devices are good at sequential writes.

Key-value stores [5, 9, 25] provide a high-performance storage service with a hash table-like interface, e.g., `GET(key)` and `PUT(key, value)`. They are typically optimized to handle a large number of small items efficiently in space, time, or both. A key-value store is a useful building

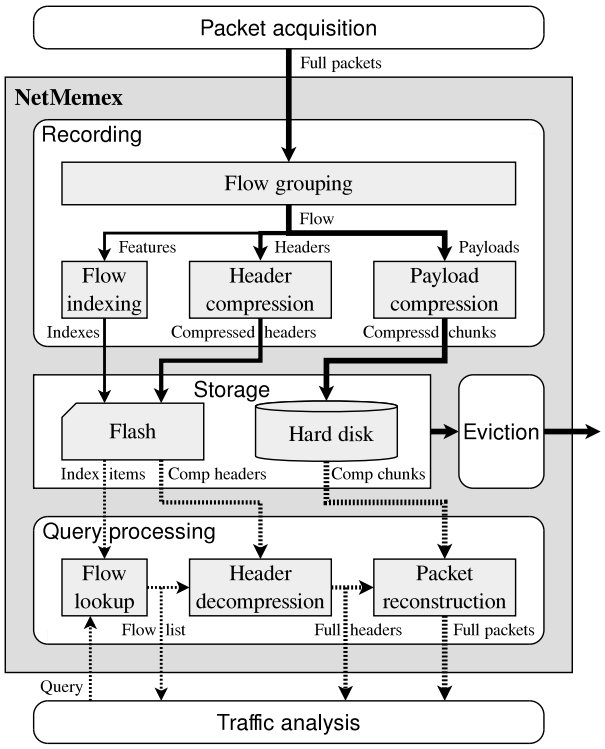

Figure 2: Architecture of NetMemex.

block for data deduplication; the system can use a chunk hash as a key to check if the chunk has been seen by looking its hash.

Flash and hard disk devices are commonly used storage devices nowadays. Flash is a NAND-based solid-state storage technology that provides high random read speed, often exceeding 35 K small reads per second [2]; however, its random write speed is slow (e.g., 300 small writes per second); flash cannot perform in-place updates (a large "erase block" must be erased before being written) and typically uses a log structure internally, which is still unable to provide high random write throughput for small writes.

Hard disks are mechanical storage devices; it uses a spinning disk and moving I/O head. Due to physical limitations, hard disks are bad at random access. On the other hand, they offer fast sequential read and write and provide inexpensive storage space; as of 2013, a hard disk drive is 15.85 times cheaper than a flash drive in terms of $ per GB.

3. DESIGN

In this section, we first provide an overview of the design of NetMemex, and then we describe each component of NetMemex in detail.

3.1 Architecture

Figure 2 depicts the architecture of NetMemex. Within the basic structure of a traffic archiving system, NetMemex combines multiple techniques to manipulate network traffic data.

Packets are acquired using packet capture software and transferred to NetMemex in the standard libpcap format.

NetMemex records the full packets by transforming them into a compressed form that can be selectively accessed. NetMemex first groups the packets into flows (§3.2). It applies header compression (§3.3) and payload compression (§3.4) to the headers and payloads of each flow, followed by flow indexing (§3.5). Through this process, NetMemex makes traffic data smaller and easy for query processing (§3.6), where each step preserves the full fidelity of the original traffic.

NetMemex stores data using both flash drives and hard disk drives, to take advantage of the strength of each storage class; flash holds headers and indexes, which are small and frequently accessed, and disk holds bulk payload data on disk.

On both storage device types, NetMemex maintains data in a log-structured way in order to handle recording and eviction efficiently (§3.7), and to extend the lifetime of the flash drives.

To enable high-speed query processing, NetMemex allows an analysis application to choose how much detail about the original traffic is necessary. For fastest operation, the application can ask NetMemex if any matching flow exists. It can retrieve a full list of matching flows if necessary. If packet headers matter, but high query performance is still needed, NetMemex can return full headers from flash drives. When the full fidelity is necessary, the application can request full packet reconstruction, which will make NetMemex fetch payload data from hard disk and construct a view of the original traffic.

Query processing produces full headers and packets in the libpcap format so that existing traffic analysis applications can easily use query results.

3.2 Flow Grouping and Lifetime

A *flow* in NetMemex is a sequence of packets with the same 5-tuple (source and destination IP address, protocol, source and destination port number if exist), ordered by arrival time (packet timestamp). *Flow grouping* is the first recording step that transforms the packet stream into a query-friendly form and facilitates per-flow compression in the subsequent recording process by grouping input packets into flows.

Figure 3 illustrates how flow grouping works. NetMemex classifies incoming packets according to their 5-tuple and enqueues them to corresponding buffers. When a 5-tuple stream terminates, NetMemex flushes all packets from the buffer of the 5-tuple, and these packets form a flow. It sends completed flows to the next recording stage to compress and index them.

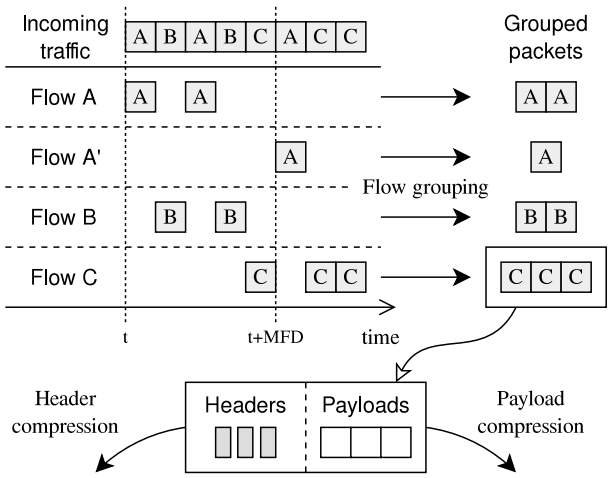

Figure 3: Example of flow grouping in NetMemex. The 5-tuple ID of each packet are labeled in the square. The packets of the 5-tuple `A` are grouped into two flows because the lifetime of any flow is disallowed from exceeding the maximum flow duration (MFD). Packets of each flow are separated into headers and payloads for the subsequent processing of header compression and payload compression.

NetMemex imposes a *maximum flow duration* (MFD) on each flow to bound the memory consumption of packet buffering. Without limits on flow duration, flow grouping could demand an unbounded amount of memory for buffers when a flow runs for a long period. To solve this problem, NetMemex stops adding new packets to a flow if its lifetime is about to exceed the MFD. In Figure 3, for instance, 5-tuple `A` actually generates two flows in NetMemex, since the flow duration would exceed if all packets are grouped into a single flow. As a consequence, NetMemex can regulate the total buffer size to be within (input traffic rate) $\times$ MFD, and the memory requirement of flow grouping becomes more predictable and manageable.

Flow grouping does not destroy any of the traffic information because each packet holds the timestamp. When required, analysis programs can sort packets by recorded timestamps to obtain the original packet arrival order.

3.3 Header Compression

In this paper, a *header* refers to a packet header and per-packet metadata (i.e., timestamp, packet lengths). Headers contain key information for queries, such as hosts, port numbers, and TCP sequence numbers.

NetMemex applies a variant of *header compression* [10, 19] to packet headers in order to efficiently reduce their size. Conventional compression algorithms, such as zlib [17] and LZO [23], aim to find and compress exact byte string matches; unfortunately, these algorithms do not work well directly for headers, as headers are a formatted data structure consisting of short and varying data fields. Instead, header compression typically encodes differences between two consecutive headers instead of recording the full header content.

To maximize the effectiveness of header compression, it should be applied to a flow consisting of packets with the same 5-tuple. This is seamlessly done in NetMemex because of the flow grouping step.

Compared to the standard header compression, NetMemex's header compression (Figure 4) is unique in two aspects:

NetMemex applies *intra-packet* compression extensively, in addition to *intra-flow* compression of the conventional header compression. NetMemex can predict header field values using the other information within the same packet and store only differences between the predicted values and the actual values in the header. For example, the IP packet length is typically inferable from the total packet length specified in the packet capture metadata. The IP checksum is also a field that can be often omitted for storage when the checksum is valid; if the checksums do not match, NetMemex simply stores the invalid checksum found in the header to preserve the full fidelity. This intra-packet compression is possible in NetMemex but not typically in the general header compression due to possible transmit errors; NetMemex do not need to worry about errors as NetMemex does not send compressed headers over network.

Because NetMemex sees multiple packets in the same flow, NetMemex can use conventional dictionary compression as the final step of header compression. This allows removing further redundancy; for instance, when TCP sequence numbers simply increase by a constant, the standard header compression would result in the same difference value for each packet; however, by applying dictionary compression, NetMemex can compress such repetition.

One (often negligible) side effect of header compression is that compressed header data becomes *flow-addressable*, not packet-addressable. One header can be read only after the decompression of the preceding header. Hence, each flow is randomly accessible efficiently, but each packet in a flow requires partial or full decompression of the flow. Since many queries request the first part of a flow or the whole flow rather than accessing random packets within it, this flow-addressability typically does not decrease query processing throughput.

After header compression, NetMemex writes the data to flash storage. The new data is appended to the most recently written one. NetMemex can easily delete the oldest headers by treating the storage as a circular queue.

3.4 Payload Compression

In packet payloads, NetMemex takes a different strategy to reduce their volume because of their redundancy characteris-

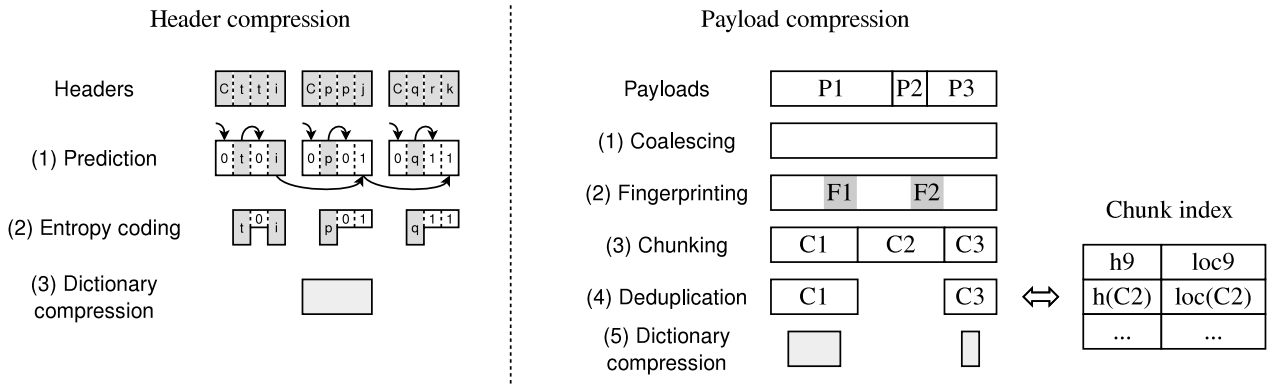

Figure 4: Schematic diagram of header compression and payload compression showing how a flow with three packets is compressed.

tics. First, the internal structure of packet payloads is more irregular than headers. Header compression largely relies on well-defined relations between packets. However, many transport layer protocols, such as TCP, break this type of redundancy in the payload data. For instance, they can adjust packet boundaries for various reasons (e.g., MSS, congestion window). This adjustment shifts and cuts the byte stream of the content, making the redundancy between two packets less explicit. Second, payloads frequently show abundant redundancy across different flows, apart from each other in time. For example, two separate connections may download the same file with a large time gap (e.g., a few hours). Because header compression is applied to each flow, it is less effective for these kinds of redundancy in the payload.

Instead, as illustrated in Figure 4, NetMemex compresses payloads using variable-length chunk deduplication, similar to LBFS [27], in conjunction with per-chunk dictionary compression. It begins by constructing a byte stream by coalescing all the payloads in a flow, which essentially ignores packet boundaries. Then, chunking is done by scanning the stream and detecting chunk boundaries using fingerprinting [32]; once chunks are discovered, each chunk is hashed using a collision-resistant hash function (e.g., SHA-1). These chunk hashes act as keys when NetMemex queries the *chunk index*, an external key-value store, to determine whether each chunk is duplicate. Finally, NetMemex stores the location of the chunks together with the compressed headers.

The chunk index should be high-performance and cost-effective. For a typical average chunk size of 4 KiB, 1 Gbps-rate traffic generates 30.5 K chunks per second; each chunk will incur one lookup for the index, so the chunk index must be fast to avoid stalling the recording process. In addition, although payload compression may reduce the volume of the payload data, this savings can be compromised if the cost of maintaining the chunk index is high. Even with the small chunk descriptor size (20-byte chunk hash and 8-byte chunk location for each chunk), the chunk index is often too large to fit entirely in main memory (e.g., 73.8 GB for one day for the above setting).

The chunk index keeps track of recent chunks only for a certain duration (*deduplication window*). If a chunk is older than this duration, its entry is removed from the chunk index (but this does not remove the chunk itself from disk); this mechanism controls how much system resources NetMemex should use for deduplication.

The above performance and space requirements make a flash-based key-value store a good candidate for the chunk index. We use SILT [22] as the chunk index in NetMemex; it provides fast operation using flash and has high space efficiency, satisfying NetMemex's requirements.

3.5 Indexing

Once a flow is compressed, NetMemex indexes it for accelerated query processing. NetMemex prepares indexes for common criteria such as IP addresses, port numbers, and so on; whenever a query includes any of such criteria, NetMemex queries the index to significantly reduce the amount of data to read and inspect.

For lightweight indexing, NetMemex's index is (1) flow-oriented, (2) compressed, (3) allows false positives, and (4) per-epoch.

NetMemex indexes flows, not individual packets. As shown in Figure 5, NetMemex builds a hash table with 65,536 buckets; it hashes the indexed field value (e.g., IP address) to determine a bucket, and inserts the location of the flow to the bucket. Since all packets in the flow are stored consecutively by header compression, NetMemex can read the packets using the flow location only.

NetMemex compresses these flow locations in each bucket using a similar way to header compression. The location numbers are monotonely increasing after sort, and thus NetMemex can store the differences between two consecutive locations. After compression, the index requires only 6.07 bytes per flow per field.

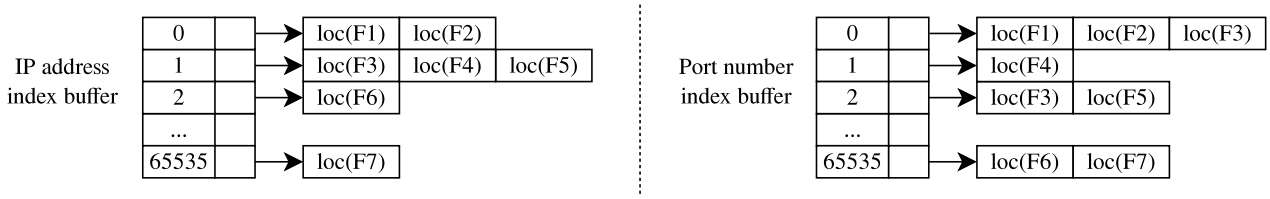

Figure 5: Lightweight flow indexing for the IP address and the port number. Each bucket contains the location of flows whose address or port number is hashed into that bucket number.

However, there can exist hash collisions from using a small number for buckets, and this leads to false positive answers (but no true negatives). NetMemex ensures that the output contains only true positive results, by decompressing the packet headers and checking if the query criteria are met.

Finally, NetMemex generates a set of independent indexes for each *epoch*. The indexes for an epoch describe the flows that have ended during that period of time. Since the epoch length is relatively short (typically a few minutes), NetMemex can construct indexes completely in memory and dump them to flash at the end of each epoch as a form of log structure. It also facilitates support for time range queries (e.g., last 1 minute) by allowing query processing to use only a subset of the indexes and makes eviction easy by removing the oldest set of indexes.

3.6 Query Processing

NetMemex supports two query modes: offline and online. In offline mode, NetMemex handles queries while not recording any incoming traffic. This mode is useful when rapid and quick query processing is crucial rather than accepting new packets. On the other hand, when online, NetMemex records new traffic data and processes queries using idle CPU cycles and I/O. Online mode is useful when the user needs to query the archived data without disrupting recording.

Query processing in NetMemex is made efficient by using indexes and data organization done in the recording stage. A query handling task consists of flow lookup, header decompression, and packet reconstruction. Each step can generate a list of flows, full headers, and full packets, respectively, based on the result from the previous step. Depending on the query type, NetMemex determines how further query processing it should proceed.

Flow lookup is the first stage of query processing. This step uses indexes that can help refining query hits. When given a time range in the query, NetMemex chooses a subset of the indexes stored; NetMemex looks up selected indexes and intersects the lookup results (i.e., a set of flow locations) if there are any *AND* operation. Then, it verifies the full index keys stored together with compressed headers to finally refine the flow list.

Flow lookup is useful in answering simple and quick queries. For example, existence test queries, which asks whether there is any flow seen in the traffic, can be done by using the flow lookup only.

The header decompression stage adds header content to the list of matching flows. NetMemex proceeds by decompressing the stored headers and emitting libpcap-compatible records.

The last query processing step is packet reconstruction, which builds and outputs full packets. This stage uses the list of chunk locations attached to the compressed headers. The specified chunks are read from disk, decompressed, and combined as a byte stream for a flow. Then, using the payload length information in the headers, the byte stream is divided into packet payloads and appended to each header. This stage is slower to process than the previous steps because reading chunks requires hard disk access, and decompressing chunks may cost a large amount of computation compared to the previous stages.

3.7 Eviction

Eviction is simple in NetMemex because it stores all traffic data in a log-structured way. It can remove the oldest data without affecting the newer data.

NetMemex currently does not send the evicted traffic data to an external longer-term archival service, which remains as future work.

4. IMPLEMENTATION

Our implementation of NetMemex uses multiple CPU cores to provide high-speed recording and query processing while avoiding packet drops.

The main thread of NetMemex accepts input packets from either a file or from standard input in the libpcap [37] trace format, which allows flexibility in choosing the packet acquisition method. This thread enqueues input packets to a buffer corresponding to their 5-tuple. When it detects the end of a flow, it adds index entries for the flow in the flow indexes; then, it dequeues and dispatches the flow to a worker thread.

The worker thread compresses the header and payload of the packets in the flow and stores the result to flash and disk.

The main thread also handles queries in online and offline mode. Query processing uses continuations to implement online mode. NetMemex can stop query processing when it needs to continue archiving; it can resume the query processing task from the last state where it left without losing the work done before.

NetMemex is implemented in 17 K lines of C++ and is targeted to run on GNU/Linux x86-64. For dictionary compression in both header and payload compression, NetMemex uses LZO [23], which makes a good compromise between compression ratios, compression speed, and decompression speed.

5. EVALUATION

In this section, we evaluate three aspects of NetMemex: (1) query performance; (2) recording throughput and memory use; and (3) storage use.

5.1 Methodology

Evaluating a traffic archiving system often requires full-packet network traffic. This imposes two practical issues in privacy concerns and fair experimental comparisons.

Access to actual packet data must be handled carefully. Both we and our IRB required that we avoid storing any traffic data possibly containing personally identifiable information (PII) on non-volatile storage, even if it is encrypted with a known encryption key. This restriction means that payload processing had to be performed on the live data on-the-fly.

However, this live payload processing makes it difficult to fairly compare multiple system configurations. Different experiment runs with a certain system setting will see different live traffic data, which may lead to completely different results. If one experiment run uses multiple system configurations at a time, the system must handle more tasks than would have been required in a realistic situation with one system configuration.

To mitigate these difficulties, we investigate both *static* and *dynamic* aspects of NetMemex's performance. In particular, we used the original live traffic with full payloads to investigate storage use under different system settings, and we also used synthetic traffic based on anonymized header-only traces to examine system throughput and memory and storage use.

The following table shows the hardware we used for the evaluation.

Component	Specification
CPU	2x Intel Xeon E5-2450 at 2.1 GHz
DRAM	DDR3 PC3-12800 48 GiB
HDD (input files)	Western Digital RE4 500 GB
HDD (output files)	2x Seagate Barracuda LP 2 TB (RAID 0)
SSD (output files)	Intel X25-M G2 80 GB

5.2 Input Workloads

We used two types of input workloads, `UNIV*` and `ISP*`, as summarized in Table 1.

For `UNIV*` workloads, NetMemex directly reads live traffic from a border router of a university. Due to privacy concerns, NetMemex performed the entire archival process on the live traffic without storing data with any PII on non-volatile storage devices.[1]

To measure the efficiency of payload compression, NetMemex chunked, hashed, and compressed the traffic data using multiple chunking configurations (chunk size, variable- or fixed-length chunks). It recorded the hashes and both original and compressed sizes. Because performing this extra chunking and compression required significant system resources, we restricted NetMemex to process only the portion of the university border traffic that originates from or is destined for eight randomly sampled /24 subnets within the university network. The specific identities of the subnets were kept anonymous from us. Because traffic subsampling typically reduces redundancy within the input, we believe that this measurement method makes our estimates of space savings from deduplication conservative.

The `ISP*` workload uses the CAIDA Anonymized Internet Traces 2012 Dataset [1]. These traces contain headers only (i.e., timestamps, packet sizes, and anonymized packet headers).

This information is insufficient to provide enough data to evaluate NetMemex because of its lack of payloads. Therefore, we appended *synthetic payloads* to each packet for two main purposes: (1) as we can control the type of payload generation, we can evaluate NetMemex from different angles with high reproducibility, which is often impossible with real traffic data; and (2) by feeding artificial payload data into NetMemex, we can evaluate NetMemex's behavior under uncommon but important workloads, for example, anomalous network traffic or attacks targeted at NetMemex itself.

In `ISP*` workloads, we used two types of payloads: redundant (`-RE`) and non-redundant (`-NR`); redundant payloads set each *i*-th payload byte to a product of *i* and a random value specific to the packet, whereas non-redundant payloads use completely random bytes for the payload. `ISP-RE` shows how NetMemex would behave with highly redundant traffic, and `ISP-NR` often incurs the heaviest load on NetMemex because it cannot reduce the amount of the total data it must handle through data compression.

Table 2 shows common system parameters used for each input workload. We used a longer maximum flow length for `UNIV*` because their relatively low bandwidth allows NetMemex to store long flows for flow grouping, whereas we limit it to 1 minute under `ISP*` workloads to avoid running out of system memory.

[1]The traffic collection and analysis was conducted with IRB approval for case # *anonymized for submission*.

Source	Description	Start Time (UTC)	Length	Original Bitrate	Packet Count	Avg Packet Size
UNIV1	University border router	2011-07-28 19:17	7 days	5.58 Mbps	980 M	430 B
UNIV2	University border router	2011-08-23 17:26	7 days	4.08 Mbps	739 M	418 B
ISP*	CAIDA equinix-sanjose B	2012-11-15 13:00	30 minutes	1.57 Gbps	527 M	669 B

Table 1: Input workloads. The packet size excludes metadata (e.g., timestamps).

Input Workload	Epoch Length	Maximum Flow Length
UNIV1	1 minute	5 minutes
UNIV2	1 minute	5 minutes
ISP-RE	1 minute	1 minute
ISP-NR	1 minute	1 minute

Table 2: The epoch length and the maximum flow length used in the experiments.

All results used variable-length chunking with a target chunking size of 4096 bytes, if not specified; UNIV* additionally used variable-length chunking with target average chunk size of 256 and 1024 bytes, and fixed-length chunking with target chunk size of 1024 and 4096 bytes, and no-chunking (treats the whole flow's payload as one chunk).

5.3 Query Performance

To demonstrate query performance, we used ISP-NR. We omit the ISP-RE result because ISP-NR is more challenging traffic to record due to its data volume and thus we can expect conservative performance results. All query outputs were written to /dev/null to prevent the measurement from being affected by the query output size.

Query types Query performance greatly varies by query type, e.g., match criteria, output data formats, query mode, and so on, as it performs different data structure access and computation to serve the queries.

We combine the following four characteristics to generate various query types in the query performance evaluation:

1. *Query mode* controls whether NetMemex should stop recording while processing queries ("offline") or should continue to record new traffic while handling query processing using idle system resources ("online").
2. *Time range* defines what period of time NetMemex should search for. "Last 1 min" queries the traffic data recorded in the last 1 minute, while "entire" looks for all recorded traffic data.
3. *Retrieval format* specifies how much information the query wants. "Existence test" queries only ask if there exists any flow or packet that matches query criteria. "Header" and "full" queries request full headers and payloads for the matching packets.
4. *Query target* indicates how frequent hit the query will cause. "No hits" means there will be no result at all matching the query criteria (e.g., `source address == 1.1.1.1`). "light hitters" will have a few results in the recorded traffic (e.g., infrequently used ports), whereas "heavy hitters" will result in a large number of flows and packets as the query result (e.g., HTTP connections).

Offline query mode Table 3 shows NetMemex's performance in offline query mode. With the query time range limited to the last 1 minute, even the slowest existence test query, which has no hits, thus making NetMemex enumerate all candidate matches returned by the flow index, can be processed in 38.4 µs per query on average. The query time jumps to 2.74 ms per query when NetMemex is instructed to investigate the entire time range. Other query results involving light hitters or heavy hitters return in at most 15.1 µs on average because NetMemex often finds an early match and can stop further query processing.

NetMemex can return a high number of flows and packets for queries requesting headers. It can handle 24.5 K flows per second, or 537.7 K packets per second at minimum. This high speed is due to the fact that NetMemex can find matching packets quickly using flow indexes while making I/O to the fast flash drive only.

For retrieving full packets with payloads, NetMemex shows about 0.1–5.8 K flows per second, or 1.6–98.4 K packets per second because NetMemex accesses disk to read chunks to reconstruct full packet data.

In retrievals involving header or payload retrieval, heavy hitters tend to allow higher query performance because NetMemex spends less time enumerating unused candidate flows.

When indexes cannot help query processing (e.g., using a TCP sequence number with no other query criteria), NetMemex examined 97.3 K flows per second, or 1,670.7 K packets per second, similar to heavy hitters' throughput.

Comparison to query processing on pcap-format files To demonstrate how NetMemex can make query processing interactive, we performed existence test queries on ISP* *header-only* pcap files with no compression, zlib, or LZO. For the no-hit query type, queries on pcap required 179.4, 257.5, 174.2 **seconds** per query, respectively for each compression method; tcpdump's raw packet header scanning speed was at most 177.4 K flows per second, or 3,024.3 K packets per second, which was faster than NetMemex's exhaustive packet header scanning without using indexes, but when NetMemex

Time Range	Query Target	µsec/query
Last 1 min	No hits	38.4
Last 1 min	Light hitters	12.7
Last 1 min	Heavy hitters	11.1
Entire	No hits	2,740.9
Entire	Light hitters	15.1
Entire	Heavy hitters	13.2

Retrieval	Time Range	Query Target	Flows/sec	Packets/sec
Header	Last 1 min	Light hitters	74,710.3	1,487,702.6
Header	Last 1 min	Heavy hitters	99,322.6	1,594,282.1
Header	Entire	Light hitters	24,450.9	537,654.5
Header	Entire	Heavy hitters	84,704.6	1,563,395.2
Full	Last 1 min	Light hitters	112.6	1,601.0
Full	Last 1 min	Heavy hitters	5,819.3	98,434.1
Full	Entire	Light hitters	93.6	1,870.3
Full	Entire	Heavy hitters	5,814.9	98,384.2

Table 3: Query performance in offline mode on `ISP-NR`. The left table shows the existence test results, the right table shows the packet retrieval results.

Time Range	Query Target	µsec/query
Last 1 min	No hits	87.9
Last 1 min	Light hitters	42.4
Last 1 min	Heavy hitters	37.0
Entire	No hits	33,463.3
Entire	Light hitters	41.7
Entire	Heavy hitters	41.5

Retrieval	Time Range	Query Target	Flows/sec	Packets/sec
Header	Last 1 min	Light hitters	19,053.4	441,648.4
Header	Last 1 min	Heavy hitters	30,474.5	537,329.4
Header	Entire	Light hitters	4,401.8	89,276.5
Header	Entire	Heavy hitters	24,529.9	477,144.5
Full	Last 1 min	Light hitters	10.7	177.2
Full	Last 1 min	Heavy hitters	404.3	8,732.0
Full	Entire	Light hitters	10.0	248.9
Full	Entire	Heavy hitters	1,768.5	30,917.9

Table 4: Query performance in online mode on `ISP-NR`. The left table shows the existence test results, the right table shows the packet retrieval results.

can use indexes, tcpdump is more than 4 orders of magnitude slower than the worst case for NetMemex requiring 2.74 ms per query. We expect that the performance gap will be even larger if tcpdump operates on full-packet traces because tcpdump must read full payloads to access packet headers.

This slow query processing with pcap files makes this approach inadequate in the situations where quick query processing is crucial.

Online query mode Table 4 shows the query performance for the same types of queries when NetMemex is actively recording new traffic data. The average recording speed of NetMemex is adjusted to 0.5 Gbps so that NetMemex has some idle time to process queries.

As the recording activity creates flash and disk activity, query throughput is slower during online operating than when offline. However, most existence test queries finish in less than 100 µs on average, with the exception of no-hits queries over the entire time range, which took 33.5 ms per query. The reason why only this type of query exhibits high latency is because the epoch data and flow index are less likely to be cached in the in-memory system page cache because of the new data generated by recording; NetMemex must access indexes and header data on flash, leading to relatively longer query time. Nevertheless, NetMemex handles existence test queries at a high enough rate for users to make queries interactively.

If indexes are not used, NetMemex could inspect 25.7 K flows per second, or 471.0 K packets per second, when also recording traffic.

Input Workload	Chunk Deduplication	Dictionary Compression	Throughput (Mbps)
`ISP-RE`	No	No	665.2
`ISP-RE`	Yes	No	700.4
`ISP-RE`	No	Yes	**1,153.4**
`ISP-RE`	Yes	Yes	955.1
`ISP-NR`	No	No	**664.4**
`ISP-NR`	Yes	No	629.4
`ISP-NR`	No	Yes	656.6
`ISP-NR`	Yes	Yes	619.1

Table 5: Recording performance with different payload processing modes.

5.4 Recording Performance

Table 5 presents recording performance with `ISP-RE` and `ISP-NR`. To show system components' contribution to the recording performance, we turn on and off chunk deduplication and dictionary compression in payload compression as noted in the table. We do not report the performance with `UNIV*` because we used multiple chunk size configurations simultaneously on the live traffic to avoid storing the original network traffic, and thus its performance was not representative.

Input Workload	Original Size (GB)	Compressed Size (GB) NetMemex	zlib	LZO
UNIV1	48.3	**5.13 (-89.4%)**	–	–
UNIV2	36.6	**7.26 (-80.2%)**	–	–
ISP-RE	26.9	**7.41 (-72.5%)**	14.2 (-47.2%)	16.6 (-38.3%)
ISP-NR	26.9	**7.50 (-72.1%)**	14.2 (-47.2%)	16.6 (-38.3%)

Table 6: Storage use by headers. The numbers in the parenthesis show the difference of the new chunk count and size from the original count and size.

With redundant payloads (ISP-RE), disabling chunk deduplication and enabling dictionary compression gives the best recording throughput of 1,153.4 Mbps. As we will see in Table 7, ISP-RE greatly benefits from dictionary compression in reducing the data volume, and this leads to faster operation with less I/O.

When payloads contain no redundancy (ISP-NR), using dictionary compression and chunk deduplication only adds overhead as they require additional computation and I/O.

With chunk deduplication and dictionary compression enabled, recording throughput is at most 17.2% slower than the best throughput configuration.

Thus, the result with synthetic workloads suggests that NetMemex can achieve high recording throughput with payload compression, while it can further increase recording speed by adaptively bypassing chunk deduplication when its saving turns out to be less significant for the current traffic feed (e.g., anomalous or heavy attack traffic).

Throughout recording, the amount of DRAM consumed by NetMemex did not exceed 15.7 GB.

5.5 Storage Space Use

Header As shown in Table 6, NetMemex effectively shrinks metadata and packet headers for all workloads. It saves at least 80.2% of space in header information size for UNIV1 and UNIV2, while it achieves 72.1–72.5% savings for the ISP* workloads; this savings greatly exceeds the savings from simply compressing pcap-format traces with zlib and LZO, which required 14.2 GB and 16.6 GB, respectively, for ISP*.

The university trace headers are compressed more than the ISP inputs primarily because UNIV* used longer flow length (5 minutes) than ISP* did (1 minute). By having longer flows, NetMemex's header compression becomes more effective, by not repeating a first full packet for each flow and having better efficiency for dictionary compression on the longer header data.

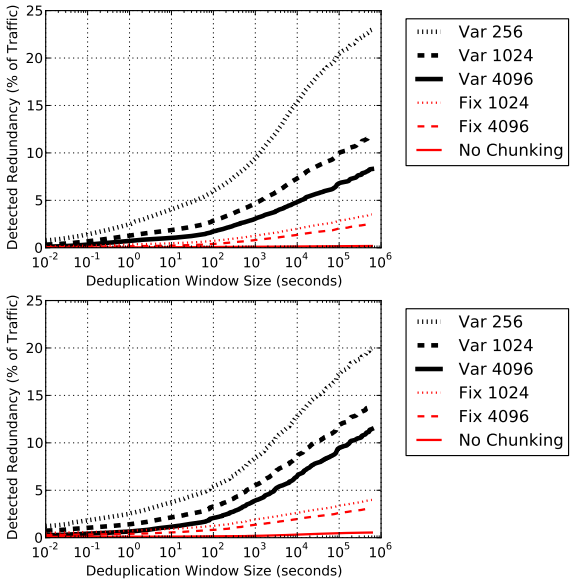

Figure 6: Detected redundancy in payload from UNIV1 (upper) and UNIV2 (lower) when varying the deduplication window. No per-chunk compression is applied.

In addition, the university has fewer hosts talking to other hosts. Common IP addresses thus appear more frequently, increasing the header compression efficiency.

Note that ISP-RE and ISP-NR have slightly different header sizes on storage because this header information includes the pointer to the payload chunk.

Payload Table 7 shows the statistics for the original and unique chunks obtained by applying payload chunking and per-chunk dictionary compression.

For redundancy elimination for payloads, NetMemex saves 8.32–11.5% for real world data in UNIV1 and UNIV2. After compression, the total savings increase to at most 17.4–19.0%.

While the compression ratios achieved for the synthetic payload ISP traces are obviously artificial, we discuss them briefly so that their effect on recording and query performance is clear. In the redundant trace, ISP-RE, chunk deduplication reduces the number of chunks by 39.9%, but only saves 7.15% of total space: Deduplication was most effective for small chunks. Dictionary compression, on the other hand, saved 84.4% of total space on this highly-compressible workload.

Optimizing payload compression for the best storage cost The storage used by payload data differs by how NetMemex chunks and deduplicates payloads. We vary the deduplication window size (how long the chunk index retains the stored chunk hash) with different chunking methods and examine the amount of redundancy NetMemex detects and removes.

Input Workload	Original Chunk Count	Original Chunk Size	Unique Chunk Count	Unique Chunk Size	Unique Chunk Compressed Size	Best-Cost Dedup Win Compressed Size
UNIV1	196.2 M	371.5 GB	184.9 M (-5.76%)	340.6 GB (-8.32%)	306.9 GB (-17.4%)	316.2 GB (-14.9%)
UNIV2	127.0 M	270.2 GB	116.3 M (-8.43%)	239.2 GB (-11.5%)	218.8 GB (-19.0%)	227.9 GB (-15.7%)
ISP-RE	57.9 M	331.6 GB	34.8 M (-39.9%)	307.9 GB (-7.15%)	51.8 GB (-84.4%)	–
ISP-NR	92.9 M	331.6 GB	92.9 M (-0.0%)	331.6 GB (-0.0%)	333.7 GB (+0.63%)	–

Table 7: Payload compression results. The numbers in the parenthesis show the difference of the new chunk count and size from the original count and size.

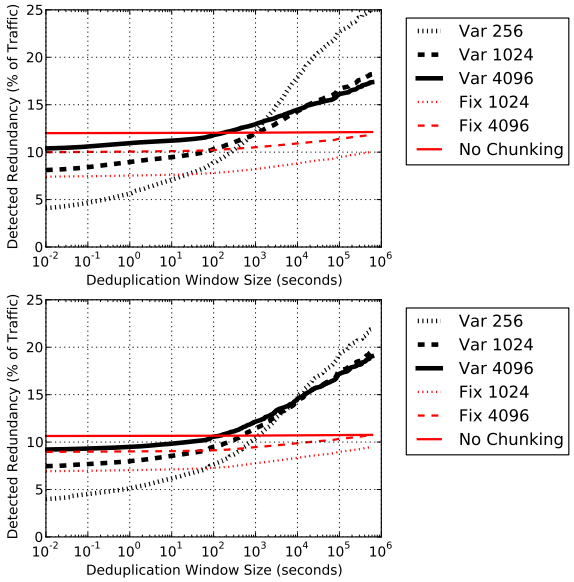

Figure 7: Detected redundancy in payload from UNIV1 (upper) and UNIV2 (lower) when varying the deduplication window. Per-chunk compression is applied.

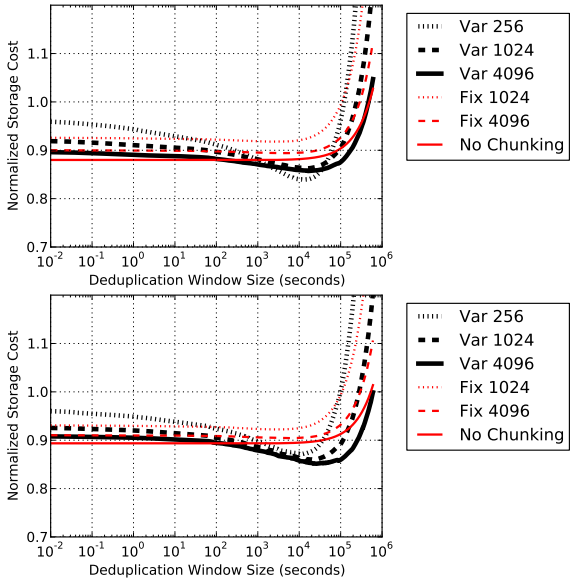

Figure 8: Storage cost analysis for storing the payload data of UNIV1 (upper) and UNIV2 (lower) when varying the deduplication window.

Figure 6 and 7 show how these factors affect the final space reduction, in uncompressed size (before applying dictionary compression) and compressed size (after applying dictionary compression; the final space use on storage).

Variable-length chunking allows more space savings than fixed-length chunking. Except for a small deduplication window (<100 seconds) where fixed-length chunks can be compressed well with dictionary compression, variable-length chunking helps find more duplicate chunks and results in better payload compression.

A larger deduplication window allows detecting a larger amount of redundant data. However, the returns diminish: beyond 10 K seconds, the slope of the detected redundancy decreases.

This diminishing effect leads to a sweet spot in storage cost savings. If NetMemex uses a large deduplication window, it saves HDD space, while using more SSD space to store the larger chunk index with more entries, and vice versa.

Figure 8 plots how storage cost is affected by applying deduplication. For this plot, we set the HDD price to $0.0467 per GB and the SSD price to $0.740 per GB, based on the market price of HDD and SSD as of January 2013. The figure clearly shows that using variable-length chunking with a small target chunk size (e.g., 256 bytes) can give the best cost effectiveness, while fixed-length chunking and no-chunking showed low redundancy.

However, a smaller chunk size has an extra cost: it makes many queries to the chunk index; thus, the chunk size should be just small enough to fully utilize the flash drive where the chunk index is stored.

The optimal deduplication window size typically lies between 10 K and 30 K seconds; a larger window increases the storage cost significantly due to too large chunk index size on flash. Using the optimal deduplication window size for variable-length chunking with the 4096-byte average chunk size target, NetMemex reduces storage cost by up to 14.9% (UNIV1) and 15.7% (UNIV2).

Index uses an insignificant amount of flush storage. For ISP* workloads, NetMemex used 199.3 MB for the IP ad-

Input Workload	Archival Method	On-flash Size (GB)			On-disk Size (GB)		Total Cost	Average of Two
		Header	Flow Index	Chunk Index	Header	Payload		
UNIV1	Original dump	–	–	–	48.3	371.5	$19.60	$16.97
UNIV2	Original dump	–	–	–	36.6	270.2	$14.33	
UNIV1	LZO-compressed dump	–	–	–	29.8*	326.9*	$16.66	$14.52
UNIV2	LZO-compressed dump	–	–	–	22.6*	242.4*	$12.38	(-14.4%)
UNIV1	NetMemex	5.13	1.59*	0.141*	–	316.2	$19.84	$18.34
UNIV2	NetMemex	7.26	0.976*	0.139*	–	227.9	$16.84	(+8.1%)

Table 8: Estimated total storage cost to archive 7-day traffic at UNIV1 and UNIV2. The numbers with a star (*) are based on our predictions. The numbers in the parenthesis show the difference of the new storage cost from the storage cost with the original dump.

dress index and 176.0 MB for the port number index. This translates to 12.14 B per flow, or only 0.712 B per packet.

Total storage cost Combined with the cost to store flow indexes and compressed headers on flash, the total storage cost using NetMemex is comparable to the cost required by packet dump on hard disk, as summarized in Table 8. In this table, flow index sizes are an estimated value based on the result with ISP⋆, and chunk index sizes are chosen based on Figure 8. The header size in the compressed dump is based on our result with LZO-compressed ISP⋆ traces while the payload size in the compressed dump is approximated using the detected redundancy using no-chunking in UNIV⋆ (Figure 7), 12.0% and 10.3%, respectively. All other sizes are directly measured quantities. We do not perform similar evaluation with ISP⋆ workloads because they lack real payload data and are too short to give a meaningful result.

Summary NetMemex provides high query performance and near-Gbps recording throughput while adding only 8.1–26.3% to the storage cost.

6. RELATED WORK

Time Machine [24] is a traffic archiving system that exploits the heavy-tailed nature of flows to reduce the total data volume. It keeps the first 10–20 KB of each connection and drops the rest; this allows retaining 91–96% of all connections while saving 90% of hard disk space. They demonstrated recording about 4 days of traffic at a research institution using 2.1 TB hard disk space. NetMemex differs from Time Machine in that NetMemex performs full traffic archival without loss, which protects NetMemex from by attacks that exploit lossy archival. In addition, NetMemex exhibits high-performance query throughput, which requires only a few µs per query; in contrast, Time Machine takes 125 ms per *in-memory* query on average and may require minutes to complete disk-based queries due to its lack of efficient data layout and indexing and the use of hard disk only. These differences in the completeness of archived data and the throughput of query processing make NetMemex more useful in retrospective analysis and other security uses, but at substantially higher storage cost.

RRDtrace [30] is another lossy traffic archiving system that focuses on storing a longer period of data on fixed storage space. Different from Time Machine, RRDtrace applies more aggressive sampling on older data, preserving more detail about recent data without requiring much storage space. While RRDtrace is stronger than Time Machine against the attacks on the lossy archival because of the non-deterministic nature of its data reduction technique, it neither exploits the redundancy in the network traffic to reduce the data volume nor provides query processing to quickly access the large amount of recorded traffic data.

Taylor *et al.* [36] presents a network data storage system that shares several commonalities with NetMemex. In this work, they aggregate packets into connections and generate summary objects that describe the payloads of the connections as a list of key-value pairs, which are indexed by a set of partitioned indexes. Unlike NetMemex, however, this technique is an application protocol-specific approach because it requires interpreting payload data to generate summary objects with application-level domain knowledge. Therefore, their system leaves a possibility for a bypass by using maliciously malformed payloads, encrypted connections (e.g., SSL), or the application protocols that are not being inspected.

NetFlow [28] and other aggregation-based monitoring techniques extract and record interesting features from the network traffic. These features are typically very small compared to the original traffic volume, enabling their efficient storage for a long term. However, it is hard to define features of interest for many security applications such as retrospective analysis because these feature are simply unknown in advance. Further, since it discards the original traffic content, they are unsuitable for applications that require the original traffic for accurate evaluation of system performance and correctness.

NET-FLi [14] and pcapIndex [15] present a compact and fast indexing scheme that can be built on packet traces in the pcap format. It creates a bitmap index whose bit indicates the presence of a certain value at the corresponding location of the original trace, then it applies a bitmap compression technique

called COMPAX. NetMemex generally generates more compact indexes than these systems; NetMemex requires 0.356 B/packet to index a field for ISP traces, whereas COMPAX indexes require 3.64 B/packet on average to index a field, making NetMemex's indexes an order of magnitude smaller than COMPAX indexes; this is possible because NetMemex performs flow-oriented data reorganization and allows its indexes to contain false positives that can be easily filtered out during query processing.

RasterZip [16] is a compression technique optimized for network traffic data, which performs column-wise compression opposed to byte-wise compression of conventional dictionary compressors. While RasterZip can compress data in an agnostic way as long as the data items are well structured (e.g., a fixed number of columns), NetMemex's header compression can handle irregularities in the input data (e.g., IP only, TCP, and UDP headers). In addition, RasterZip fails to remove redundancy that can be found with in-depth domain knowledge; NetMemex's header compression detects and removes intra-packet redundancy (e.g., valid IP checksum) because NetMemex fully understands and can exploit the semantics of packet headers.

Selective Packet Paging (SPP) [31] strengthens packet capture systems from overload attacks that cause the systems to accept packets more than they can process. NetMemex is orthogonal to this work and can take advantage of it because SPP is used in the packet acquisition stage, which is outside of NetMemex.

7. CONCLUSION

NetMemex enables high-throughput recoding and query processing of network traffic without sacrificing the fidelity of the archived information or requiring high storage expenses. In this paper, we show that NetMemex can handle common types of queries within a few µs on average with high-speed flow/packet retrieval, record near-Gbps full packet traffic, and use a small amount of storage space whose total hardware cost is comparable to hard disk-only solutions. NetMemex achieves these goals simultaneously by performing data reorganizations, using flash and disk carefully, and applying smart data compression, in a sophisticated and organized manner.

REFERENCES

[1] The CAIDA UCSD Anonymized Internet Traces 2012 - 20121115. http://www.caida.org/data/passive/passive_2012_dataset.xml, 2012.

[2] Intel Solid-State Drive X25-M Series. http://ark.intel.com/products/56601/Intel-SSD-X25-M-Series-80GB-2_5in-SATA-3Gbs-34nm-MLC, 2012.

[3] B. Aggarwal, A. Akella, A. Anand, A. Balachandran, P. Chitnis, C. Muthukrishnan, R. Ramjee, and G. Varghese. EndRE: an end-system redundancy elimination service for enterprises. In *Proc. 7th USENIX NSDI*, Apr. 2010.

[4] A. Anand, V. Sekar, and A. Akella. SmartRE: an architecture for coordinated network-wide redundancy elimination. In *SIGCOMM '09: Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, pages 87–98. ACM, 2009.

[5] D. G. Andersen, J. Franklin, M. Kaminsky, A. Phanishayee, L. Tan, and V. Vasudevan. FAWN: A fast array of wimpy nodes. In *Proc. 22nd ACM Symposium on Operating Systems Principles (SOSP)*, Oct. 2009.

[6] E. Anderson and M. Arlitt. Full packet capture and offline analysis on 1 and 10 Gb/s networks. Technical Report HPL-2006-156, HP Laboratories Palo Alto, Nov. 2006.

[7] Y. Cheng, J. Bellardo, P. Benko, A. C. Snoeren, G. M. Voelker, and S. Savage. Jigsaw: Solving the puzzle of enterprise 802.11 analysis. In *Proc. ACM SIGCOMM*, Aug. 2006.

[8] C. Cranor, T. Johnson, O. Spataschek, and V. Shkapenyuk. Gigascope: a stream database for network applications. In *Proc. ACM SIGMOD*, pages 647–651, June 2003.

[9] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels. Dynamo: Amazon's highly available key-value store. In *Proc. 21st ACM Symposium on Operating Systems Principles (SOSP)*, Oct. 2007.

[10] M. Degermark, B. Nordgren, and S. Pink. *IP header compression*. Internet Engineering Task Force, Feb. 1999. RFC 2507.

[11] L. Degioanni and G. Varenni. Introducing scalability in network measurement: toward 10 Gbps with commodity hardware. In *IMC '04: Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, pages 233–238. ACM, 2004.

[12] P. J. Desnoyers and P. Shenoy. Hyperion: high volume stream archival for retrospective querying. In *ATC'07: 2007 USENIX Annual Technical Conference on Proceedings of the USENIX Annual Technical Conference*, pages 1–14. USENIX Association, 2007.

[13] Endace Limited. http://www.endace.com/high-speed-packet-capture-hardware.html, 2013.

[14] F. Fusco, M. P. Stoecklin, and M. Vlachos. NET-FLi: on-the-fly compression, archiving and indexing of streaming network traffic. *Proc. VLDB Endow.*, 3(1-2), Sept. 2010.

[15] F. Fusco, X. Dimitropoulos, M. Vlachos, and L. Deri. pcapIndex: an index for network packet traces with legacy compatibility. *SIGCOMM Comput. Commun. Rev.*, 42(1):47–53, 2012. ISSN 0146-4833.

[16] F. Fusco, M. Vlachos, and X. Dimitropoulos. RasterZip: Compressing network monitoring data with support for partial decompression. In *Proceedings of the 12nd ACM SIGCOMM conference on Internet measurement*, IMC '12, 2012.

[17] J.-L. Gailly and M. Adler. zlib. http://www.zlib.net/, 2013.

[18] R. Geambasu, T. Bragin, J. Jung, and M. Balazinska. On-demand view materialization and indexing for network forensic analysis. In *NETDB'07: Proceedings of the 3rd USENIX*

international workshop on Networking meets databases, pages 1–7. USENIX Association, 2007.

[19] V. Jacobson. *Compressing TCP/IP Headers for Low-speed Serial Links*. Internet Engineering Task Force, Feb. 1990. RFC 1144.

[20] A. Kawaguchi, S. Nishioka, and H. Motoda. A flash-memory based file system. In *Proc. USENIX Annual Technical Conference*, Jan. 1995.

[21] J. Lilley, J. Yang, H. Balakrishnan, and S. Seshan. A Unified Header Compression Framework for Low-Bandwidth Links. In *Proc. ACM Mobicom*, Aug. 2000.

[22] H. Lim, B. Fan, D. G. Andersen, and M. Kaminsky. SILT: A memory-efficient, high-performance key-value store. In *Proc. 23rd ACM Symposium on Operating Systems Principles (SOSP)*, Oct. 2011.

[23] LZO. http://www.oberhumer.com/opensource/lzo/, 2013.

[24] G. Maier, R. Sommer, H. Dreger, A. Feldmann, V. Paxson, and F. Schneider. Enriching network security analysis with time travel. In *Proc. ACM SIGCOMM*, Aug. 2008.

[25] A distributed memory object caching system. http://memcached.org/, 2011.

[26] D. Moore, G. Voelker, and S. Savage. Inferring Internet denial of service activity. In *Proc. 10th USENIX Security Symposium*, Aug. 2001.

[27] A. Muthitacharoen, B. Chen, and D. Mazieres. A low-bandwidth network file system. In *Proc. 18th ACM Symposium on Operating Systems Principles (SOSP)*, Oct. 2001.

[28] Cisco NetFlow. http://www.cisco.com/en/US/products/ps6601/products_ios_protocol_group_home.html, 2013.

[29] Fluke Corporation. http://www.flukenetworks.com/enterprise-network/network-monitoring/Network-Time-Machine, 2013.

[30] A. Papadogiannakis, M. Polychronakis, and E. P. Markatos. RRDtrace: Long-term raw network traffic recording using fixed-size storage. In *Proceedings of the 2010 IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*, MASCOTS '10, 2010.

[31] A. Papadogiannakis, M. Polychronakis, and E. Markatos. Tolerating overload attacks against packet capturing systems. In *Proc. USENIX Annual Technical Conference*, June 2012.

[32] M. O. Rabin. Fingerprinting by random polynomials. Technical Report TR-15-81, Center for Research in Computing Technology, Harvard University, 1981.

[33] M. Rosenblum and J. K. Ousterhout. The design and implementation of a log-structured file system. *ACM Transactions on Computer Systems*, 10(1):26–52, 1992.

[34] NetScout Systems, Inc. http://www.netscout.com/products/sniffer_analysis.asp, 2013.

[35] N. Spring and D. Wetherall. A protocol-independent technique for eliminating redundant network traffic. In *Proc. ACM SIGCOMM*, Sept. 2000.

[36] T. Taylor, S. Coull, F. Monrose, and J. McHugh. Toward efficient querying of compressed network payloads. In *Proc. USENIX Annual Technical Conference*, June 2012.

[37] tcpdump/libpcap. http://www.tcpdump.org/, 2013.